

The RatioLog Project: Rational Extensions of Logical Reasoning

Ulrich Furbach · Claudia Schon · Frieder Stolzenburg · Karl-Heinz Weis ·
Claus-Peter Wirth

Abstract Higher-level cognition includes logical reasoning and the ability of question answering with common sense. The RatioLog project addresses the problem of rational reasoning in deep question answering by methods from automated deduction and cognitive computing. In a first phase, we combine techniques from information retrieval and machine learning to find appropriate answer candidates from the huge amount of text in the German version of the free encyclopedia “Wikipedia”. In a second phase, an automated theorem prover tries to verify the answer candidates on the basis of their logical representations. In a third phase — because the knowledge may be incomplete and inconsistent —, we consider extensions of logical reasoning to improve the results. In this context, we work toward the application of techniques from human reasoning: We employ defeasible reasoning to compare the answers w.r.t. specificity, deontic logic, normative reasoning, and model construction. Moreover, we use integrated case-based reasoning and machine learning techniques on the basis of the semantic structure of the questions and answer candidates to learn giving the right answers.

Keywords automated deduction · case-based reasoning · common-sense reasoning · defeasible reasoning · deontic logic · question answering · specificity

1 Rational Reasoning and Question Answering

The development of formal logic played a big role in the field of automated reasoning, which led to the development of the field of artificial intelligence (AI). Applications of *automated deduction* in mathematics have been investigated from the early years on. Nowadays automated deduction techniques are successfully applied in hard- and software verification and many other areas (for an overview see [2]).

In contrast to formal logical reasoning, however, human reasoning does not strictly follow the rules of classical logic. Reasons may be incomplete knowledge, incorrect beliefs, and inconsistent norms. From the very beginning of AI research, there has been a strong emphasis on incorporating mechanisms for rationality, such as abductive or defeasible reasoning. From these efforts, as part of the field of knowledge representation, *common-sense reasoning* has emerged as a branching discipline with many applications in AI [17].

Nowadays there is a chance to join automated deduction and common-sense reasoning within the paradigm of *cognitive computing*, which allows the implementation of rational reasoning [16]. The general motivation for the development of cognitive systems is that computers can solve well-defined mathematical problems with enormous precision at a reasonably sufficient speed in practice. It remains difficult, however, to solve problems that are only vaguely outlined. One important characteristic of cognitive computing is that many different knowledge formats and many different information processing methods are used in a combined fashion. Also the amount of knowledge is huge and, even worse, it is even increasing steadily. For the logical reasoning, a similar argument holds: Different reasoning mechanisms have to be employed and combined, such as classical deduction (forward reasoning) on the one hand, and abduction or other non-monotonic reasoning mechanisms on the other hand.

Let us illustrate this with a well-known example from the literature:

U. Furbach, C. Schon, K.-H. Weis
Artificial Intelligence Research Group, Universität Koblenz-Landau,
Universitätsstr. 1, 56070 Koblenz, Germany
E-mail: {uli, schon}@uni-koblenz.de, khweis@weis-consulting.de

F. Stolzenburg, C.-P. Wirth
Automation and Computer Sciences Department, Harz University of
Applied Sciences, Friedrichstr. 57–59, 38855 Wernigerode, Germany
E-mail: fstolzenburg@hs-harz.de, wirth@logic.at




Figure 1 The LogAnswer system uses information retrieval (IR), decision tree learning in a ranking phase, reasoning, and natural language answer generation to compute answers.

1. *Tom is an emu.*
2. *Emus are birds.*
3. *Birds normally fly.*
4. *Emus do not fly.*

The question is: Can emus fly or not? Forward reasoning allows us to infer that emus are birds and hence can normally fly. This is in conflict, however, with the strict background knowledge that emus do not fly. The conflict can be solved by assuming certain knowledge as default or defeasible, which only holds normally. Hence we may conclude here that emus and therefore Tom does not fly. We will come back to this example later (namely in Section 2.2 and 2.3).

Rational reasoning must be able to deal with incomplete as well as conflicting (or even inconsistent) knowledge. Moreover, huge knowledge bases with inconsistent contents must be handled. Therefore, it seems to be a good idea to combine and thus enhance rational reasoning by information retrieval techniques, e.g. techniques from machine learning. This holds especially for the domain of deep question answering, where communication with patterns of human reasoning is desirable.

1.1 Deep Question Answering and the LogAnswer System

Typically, question answering systems, including application programs such as *Okay Google*® or Apple®’s *Siri*, communicate with the user in natural language. They accept properly formulated questions and return concise answers. These automatically generated answers are usually not extracted directly from the web, but, in addition, the system operates on an extensive (background) knowledge base, which has been derived from textual sources in advance.

LogAnswer [8, 9] is an open-domain question answering system, accessible via a web interface (www.loganswer.de) similar to that of a search engine. The knowledge used to answer the question is gained from 29.1 million natural-language sentences of a snapshot of the German Wikipedia. Furthermore, a background knowledge consisting of 12,000 logical facts and rules is used. The LogAnswer system was developed in the DFG-funded LogAnswer project, a cooperation between the groups on Intelligent Information and

Communication Systems at the FernUniversität Hagen and the AI research group at the University of Koblenz-Landau. The project aimed at the development of efficient and robust methods for logic-based question answering. The user enters a question and LogAnswer presents the five best answers from a snapshot of the German “Wikipedia”, highlighted in the context of the relevant textual sources.

Most question answering systems rely on shallow linguistic methods for answer derivation, and there is only little effort to include semantics and logical reasoning. This may make it impossible for the system to find any answers: A superficial word matching algorithm is bound to fail if the textual sources use synonyms of the words in the question. Therefore, the LogAnswer system models some form of background knowledge, and combines cognitive aspects of linguistic analysis, such as semantic nets in a logical representation, with machine learning techniques for determining the most appropriate answer candidate.

Contrary to other systems, LogAnswer uses an automated theorem prover to compute the replies, namely Hyper [3], an implementation of the hypertableaux calculus [1], extended with equality among others. It has demonstrated its strength in particular for reasoning problems with a large number of irrelevant axioms, as they are characteristic for the setting of question answering. The logical reasoning is done on the basis of a logical representation of the semantics of the entire text contained in the Wikipedia snapshot. This is computed beforehand with a system developed by computational linguists [13] which employs the MultiNet graph formalism (Multilayered Extended Semantic Networks) [15].

Since methods from natural-language processing are often confronted with flawed textual data, they strive toward robustness and speed, but often lack the ability to perform more complex inferences. By contrast, a theorem prover uses a sound calculus to derive precise proofs of a higher complexity; even minor flaws or omissions in the data, however, lead to a failure of the entire derivation process. Thus, additional techniques from machine learning, defeasible and normative reasoning etc. should be applied to improve the quality of the answers — as done in the RatioLog project.

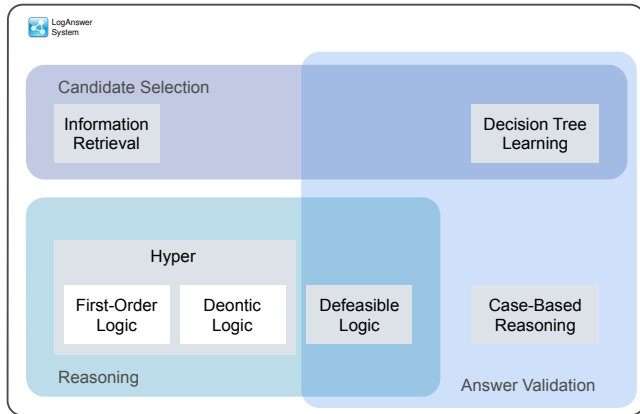


Figure 2 Techniques used in the different modules of the LogAnswer system.

For this, the reasoning in classical logic is extended by various forms of non-monotonic aspects, such as defeasible argumentation. By these extensions, the open-domain question answering system LogAnswer is turned into a system for rational question answering, which offers a testbed for the evaluation of rational reasoning.

1.2 The LogAnswer System and its Modules

When processing a question, the LogAnswer system performs several different steps. Figure 1 presents details on these steps. At first, information retrieval (IR), i.e. pattern matching, is used to filter text passages suitable for the given question from the textual representation of the Wikipedia. For this, the text sources are segmented into sentence-sized passages. The corresponding representation can be enriched by the descriptors of other sentences that result from a coreference solution of pronouns, where the referred phrase is added to the description, if e.g. the pronoun 'he' refers to the individual 'Ian Fleming'. Then decision tree learning ranks the text passages and chooses a set of answer candidates from these text passages (Ranking step in Figure 1). Here, features like the number of matching lexemes between passages and the question or the occurrences of proper names in the passage are computed. Up to 200 text passages are extracted from the knowledge base according to this ranking.

In the next step (Reasoning), the Hyper theorem prover is used to check if these text passages provide an answer to the question. For every answer candidate, a first-order logic representation of both the question and the answer candidate is combined with a huge background knowledge. These proofs provide the answer to the question by means of variable assignments. The proofs for the answer candidates are then ranked again using decision tree learning (in the answer validation phase). For the five best answers, text passages providing the answer are highlighted and presented to

the user. This is done by a natural language (NL) answer generation module, which eventually yields the final answer candidates, in our case that Ian Fleming is a British author.

In the LogAnswer system, various techniques work interlocked. See Figure 2 for an overview of the different techniques together with the modules in which they are used. Extraction of text passages for a certain question is performed in the candidate selection module. In this module, both information retrieval and decision tree learning work hand in hand to find a list of answer candidates for the current question. For each answer candidate, the reasoning module is invoked. This module consists of the Hyper theorem prover, which is used to check if the answer candidate provides an answer for the question. Since Hyper is able to handle first-order logic with equality and knowledge bases given in description logic, it is possible to incorporate background knowledge given in various (formal) languages.

An interesting extension of usual background knowledge is the use of a knowledge base containing normative statements formalized in deontic logic. These normative statements enable the system to reason in a rational way. Since deontic logic can be translated into description logics, Hyper can be used to reason on such knowledge bases. Reasoning in defeasible logic is another technique contained in the reasoning module of the LogAnswer system. With the help of defeasible logic reasoning, different proofs produced by Hyper are compared. The proofs found by Hyper provide answers to the given question by means of variable assignments. Comparing the proofs for different answer candidates therefore is used to determine the best answer. Hence defeasible logic is contained in the answer validation module. In addition to that, the answer validation module contains decision tree learning to rank different proofs found by Hyper and case-based reasoning. Details on the use of case-based reasoning and reasoning in defeasible logic, that can both be used in the answer validation phase (see Figure 1), can be found in the Section 2.

2 Searching for Good Answers

As depicted before, the reasoning component of the LogAnswer system delivers proofs, which represent the possible answers to the given question. The proofs are ranked by decision trees which take into account several attributes of the reasoning process together with the attribute from the previous information retrieval step.

In addition to this ranking, we experiment with different other techniques to improve the evaluation of answers. These are case-based reasoning (CBR) (Section 2.1), defeasible reasoning (Section 2.2), and normative (deontic) reasoning (Section 2.3). To perform systematic and extensive tests with LogAnswer, we used the CLEF database, strictly speaking, its question answering part.

CLEF stands for cross-language evaluation forum, see www.clef-campaign.org. It is an international campaign providing language data in different languages, e.g. from newspaper articles. Its workshop and competition series contains a track on question answering. We used data from CLEF-2007 and CLEF-2008 [12, 18].

2.1 CBR Similarity Measures and Machine Learning

Answer validation can be enhanced by using experience knowledge in form of cases in a case base. The resulting system module is designed as a learning system and based on a dedicated CBR control structure. Contrary to common procedures in natural-language processing, however, we do not follow the textual approach, where experiences are available in unstructured or semi-structured text form, but use a structured approach along the lines of [4]. This is possible because the knowledge source is available not only in textual but also in a logical format. The semantics of the natural-language text is given basically by first-order predicate logic formulae, which are represented by MultiNet graphs [15]. Our basis is a manually achieved classification for each pair of question (from the CLEF 2007 and 2008 data) and answer candidate (from the LogAnswer system) whether the answer candidate is a good one for the question. In order to compare and to define a similarity measure of the MultiNet graphs, we have developed a new graph similarity measure [14, 22] which improves other existing measures, e.g. [4, 6].

We measured the CBR system classification accuracy by running tests with a case base from the CLEF 2007 and 2008 data. Our overall test set had 254 very heterogeneous questions and ca. 15000 cases. For instance, in one of the evaluations, namely the user interaction simulation (see Figure 3), we examined the development of the results for a growing knowledge base. We simulated users that give reliable feedback to new, heterogeneous questions for which the LogAnswer system provides answers candidates. The test setting was to guess the classification of questions and answer candidates the system does not have in the knowledge base. The results show the increase of the classification accuracy with a growing number of correct cases in the case base. We performed a number of other evaluation experiments, e.g. 3- and 10-fold cross validations. For more information about the integrated CBR/Machine learning evaluation and test settings, please refer to [14, 22].

We further integrated case-based reasoning into the already existing answer selection techniques in LogAnswer (answer validation phase, see Figure 1). For this, the results of the CBR stage were turned into numeric features. A ranking model determined by a supervised learning-to-rank approach combined these CBR-based features with other answer selection features determined by shallow linguistic processing and logical answer validation. The final machine




Figure 3 The x-axis is the number of cases in the case base. The y-axis is the classification accuracy in percent, for correct and incorrect answer candidates, as well as the overall classification accuracy for the user interaction simulation.

learning ranker is an ensemble of ten rank-optimizing decision trees, obtained by stratified bagging, whose individual probability estimates are combined by averaging. When training the machine learning ranker on a case base optimized for perfect treatment of correct answer candidates, we get the best overall result in our tests with a mean reciprocal rank (MRR) of 0.74 (0.72 without CBR) and a correct top-ranked answer chosen in 61% (58% without CBR) of the cases. It is instructive to consider the usage of CBR features in the machine learning ranker, by inspecting all branching conditions in the generated trees and counting the frequency of occurrence of each feature in such a branching condition, since 10 bags of 10 decision trees were generated in the 10 cross-validation runs, there is a total of 100 trees to base results on [14, 22]. In total, 42.5% of all split conditions in the learned trees involve one of the CBR attributes. This further demonstrates the strong impact of CBR results on answer re-ranking.

2.2 The Specificity Criterion

More specific answer candidates are to be preferred to less specific ones, and we can compare them according to their specificity as follows. To obtain what *argument theories* call an *argument*, we form a pair of an answer candidate and its derivation. The derivation can be based on *positive-conditional rules*, generated from Hyper’s verifications and capturing the Wikipedia page of the answer candidate and the linguistic knowledge actually applied. Now we find ourselves in the setting of *defeasible reasoning* and can sort the arguments according to their specificity.

In defeasible reasoning, certain knowledge is assumed to be *defeasible*. *Strict* knowledge, however, is specified by *contingent facts* (e.g. in the emu example from Section 1, “Tom is an emu”) and *general rules* holding in all possi-

ble worlds without exception (e.g. “emus do not fly”). Strict knowledge is always preferred to knowledge depending also on *defeasible rules* (e.g. “birds *normally* fly”).

Already in 1985, David Poole had the idea to prefer more *specific* arguments in case of conflicting results as follows [19]: For any derivation of a given result, represented as a tree, consider the sets of all leaves that contribute to the applications of defeasible rules. An *activation set* is a set of literals from which all literals labeling such a set of leaves is derivable. Thereby, an activation set is sufficient to activate the defeasible parts of a derivation in the sense of a presupposition, *without using any additional contingent facts*.

One argument is now *more specific* than another one if all its activation sets are activation sets of the other one. This means that each activation set of the more specific argument (seen as the conjunction of its literals) must be more specific than an activation set of the other one. Note that the meaning of the latter usage of the word “specific” is just the *traditional common-sense concept of specificity*, according to which a criterion (here: conjunction of literals) is more specific than another one if it entails the other one.

We discovered several weaknesses of Poole’s relation, such as its non-transitivity: Contrary to what is obviously intended in [19] and “proved” in [20], Poole’s relation is not a quasi-ordering and cannot generate an ordering. We were able to cure all the discovered weaknesses by defining a *quasi-ordering* [23,24] (i.e. a reflexive and transitive binary relation), which can be seen as a correction of Poole’s relation, maintaining and clarifying Poole’s original intuition.

The intractability of Poole’s relation, known at least since 2003 [21], was attenuated by our quasi-ordering and then overcome by restricting the rules to instances that were actually used in the proofs found by Hyper, and by treating the remaining variables (if any) as constants. With these restrictions, the intractability did not show up anymore in any of the hundreds of examples we tested with our PROLOG implementation.

Running this implementation through the entire CLEF-2008 database, almost all suggested answer solutions turned out to be incomparable w.r.t. specificity, although our quasi-ordering can compare more arguments in practice than Poole’s original relation. One problem here is that we have to classify the rules of the CLEF examples as being either general or defeasible, but there is no obvious way to classify them. Another problem with the knowledge encoded in the MultiNet formalism is that it first and foremost encodes only linguistic knowledge, e.g., who is the agent of a given sentence. Only little background knowledge is available, such as on ontology. All data from the web pages, however, are represented by literals.

To employ more (defeasible) background knowledge we investigated other examples, such as the emu example from Section 1. Here, the formalization in first-order logic of the

natural-language knowledge on individuals can be achieved with the Boxer system [5,7], which is dedicated to large-scale language processing applications. These examples can be successfully treated with the specificity criterion and also with deontic logic (see subsequent section).

2.3 Making Use of Deontic Logic

Normative statements like “you ought not steal” are omnipresent in our everyday life, and humans are used to do reason with respect to them. Since norms can be helpful to model rationality, they constitute an important aspect for common-sense reasoning. This is why normative reasoning is investigated in the RatioLog project [10]. Standard deontic logic (SDL) [11] is a logic which is very suitable for the formalization of knowledge about norms. SDL corresponds to the modal logic K together with a seriality axiom. In SDL the modal operator \Box is interpreted as “it is obligatory that” and the \Diamond operator as “it is permitted that”. For example a norm like “you ought not steal” can be intuitively formalized as $\Box \neg \text{steal}$. From a model theoretic point of view, the seriality axiom contained in SDL ensures that, whenever it is obligatory that something holds, there is always an ideal world fulfilling the obligation.

In the RatioLog project, we experiment with SDL by adding normative statements into the background knowledge. The emu example from Section 1 contains the normative assertion

Birds normally fly.

which can be modeled using SDL as

$Bird \rightarrow \Box Flies$

and is added to the background knowledge. In addition to normative statements, the background knowledge furthermore contains assertions not containing any modal operators, e.g. something like the statement that all emus are birds. Formulae representing contingent facts, like the assertion

Tom is an emu.

in the emu example, are combined with the background knowledge containing information about norms. The Hyper theorem prover [3] can be used to analyze the resulting knowledge base. For example, it is possible to ask the prover if the observed world with the emu Tom fulfills the norm that birds usually are able to fly.

Within the RatioLog project both defeasible logic and deontic logic are used. There are similarities between defeasible logic and deontic logic. For example in defeasible logic there are rules which are considered to be not strict but defeasible. These defeasible rules are similar to normative statements, since norms only describe how the world ought

to be and not how it actually is. This is why we are also investigating the connection between these two logics within the RatioLog project.

3 Conclusions

Deep question answering does not only require pattern matching and indexing techniques, but also rational reasoning. This has been investigated within the RatioLog project as demonstrated in this article. Techniques from machine learning with similarity measures and case-based reasoning, defeasible reasoning with (a revision of) the specificity criterion, and normative reasoning with deontic logic help to select good answer candidates. If the background knowledge, however, mainly encodes linguistic knowledge — without general common-sense world knowledge — then the effect on finding good answer candidates is low. Therefore, future work will concentrate on employing even more background world knowledge (e.g. from ontology databases), so that rational reasoning can be exploited more effectively when applied to this concrete knowledge.

Acknowledgements The authors gratefully acknowledge the support of the DFG under the grants FU 263/15-1 and STO 421/5-1 *Ratiolog*.

References

1. Baumgartner, P., Furbach, U., Niemelä, I.: Hyper tableaux. In: J.J. Alferes, L.M. Pereira, E. Orłowska (eds.) Proceedings of 5th European Workshop on Logics in AI – JELIA'96, LNCS 1126, pp. 1–17. Springer (1996)
2. Beckert, B., Hähnle, R.: Reasoning and verification: State of the art and current trends. *IEEE Intelligent Systems* **29**(1), 20–29 (2014)
3. Bender, M., Pelzer, B., Schon, C.: System description: E-KRHyper 1.4 – extensions for unique names and description logic. In: M.P. Bonacina (ed.) CADE-24, LNCS 7898, pp. 126–134. Springer (2013)
4. Bergmann, R.: Experience Management – Foundations, Development Methodology and Internet-Based Applications. Springer, Berlin, Heidelberg, New York (2002)
5. Bos, J.: Towards wide-coverage semantic interpretation. In: Proceedings of 6th International Workshop on Computational Semantics IWCS-6, pp. 42–53. Tilburg, Netherlands (2005)
6. Bunke, H., Messmer, B.T.: Similarity measures for structured representations. In: Topics in Case-Based Reasoning: 1st European Workshop EWCBR-93 (1993)
7. Curran, J.R., Clark, S., Bos, J.: Linguistically motivated large-scale NLP with C&C and Boxer. In: Proceedings of the ACL 2007 Demo and Poster Sessions, pp. 33–36. Prague, Czech Republic (2007)
8. Furbach, U., Glöckner, I., Helbig, H., Pelzer, B.: Logic-based question answering. *KI – Künstliche Intelligenz* **24**(1), 51–55 (2010). Special Issue on Automated Deduction
9. Furbach, U., Glöckner, I., Pelzer, B.: An application of automated reasoning in natural language question answering. *AI Communications* **23**(2-3), 241–265 (2010)
10. Furbach, U., Schon, C., Stolzenburg, F.: Automated reasoning in deontic logic. In: M.N. Murty, X. He, R.R. Chillarige, P. Weng (eds.) Proceedings of MIWAI 2014: Multi-Disciplinary International Workshop on Artificial Intelligence, LNAI 8875, pp. 57–68. Springer (2014)
11. Gabbay, D., Horty, J., Parent, X., van der Meyden, R., van der Torre, L. (eds.): Handbook of Deontic Logic and Normative Systems. College Publications (2013)
12. Giampiccolo, D., Forner, P., Herrera, J., Peñas, A., Ayache, C., Forascu, C., Jijkoun, V., Osenova, P., Rocha, P., Sacaleanu, B., et al.: Overview of the CLEF 2007 multilingual question answering track. In: Advances in Multilingual and Multimodal Information Retrieval, pp. 200–236. Springer (2008)
13. Glöckner, I., Hartrumpf, S., Leveling, J.: Logical validation, answer merging and witness selection: A study in multi-stream question answering. In: Proceedings of RIAO-07. Pittsburgh (2007)
14. Glöckner, I., Weis, K.H.: An integrated machine learning and case-based reasoning approach to answer validation. In: Proceedings ICMLA 2012. Boca Raton (FL) (2012)
15. Helbig, H.: Wissensverarbeitung und die Semantik der natürlichen Sprache, 2nd edn. Springer, Berlin, Heidelberg, New York (2008)
16. Kelly III, J.E., Hamm, S.: Smart Machines: IBM's Watson and the Era of Cognitive Computing. Columbia Business School Publishing (2013)
17. Mueller, E.T.: Commonsense Reasoning, 2nd edn. Morgan Kaufmann, San Francisco (2014)
18. Peters, C.: Cross-language evaluation forum – CLEF 2008. *D-Lib Magazine* **14**(11/12) (2008). URL <http://www.dlib.org/dlib/november08/peters/11peters.html>
19. Poole, D.L.: On the comparison of theories: Preferring the most specific explanation. In: A. Joshi (ed.) Proc. 9th Int. Joint Conf. on Artificial Intelligence (IJCAI), 1985, Aug. 18–25, Los Altos (CA), pp. 144–147. Morgan Kaufmann (Elsevier) (1985)
20. Simari, G.R., Loui, R.P.: A mathematical treatment of defeasible reasoning and its implementation. *Artificial Intelligence* **53**, 125–157 (1992)
21. Stolzenburg, F., García, A.J., Chesñevar, C.I., Simari, G.R.: Computing generalized specificity. *Journal of Applied Non-Classical Logics* **13**, 87–113 (2003)
22. Weis, K.H.: A case based reasoning approach for answer reranking in question answering. In: M. Horbach (ed.) Informatik 2013 – Proceedings, no. 220 in GI-Edition, Lecture Notes in Informatics, pp. 93–104. Koblenz (2013). Also available at <http://arxiv.org/abs/1503.02917>
23. Wirth, C.P., Stolzenburg, F.: David Poole's specificity revised. In: C. Baral, G.D. Giacomo, T. Eiter (eds.) Int. Conf. on Principles of Knowledge Representation and Reasoning, pp. 168–177. AAAI Press (2014). Extended version available as SEKI-Report SR-2013-01 at <http://arxiv.org/abs/1308.4943>
24. Wirth, C.P., Stolzenburg, F.: A series of revisions of David Poole's specificity. *Annals of Mathematics and Artificial Intelligence* (2015). Accepted for publication



Ulrich Furbach is a Senior Research Professor of Artificial Intelligence at the University of Koblenz. His research interests include knowledge management, automated reasoning, multiagent systems, and cognitive science.



Claudia Schon is employed at the University of Koblenz-Landau in the RatioLog project. Her research interests include logic, especially description logics, artificial intelligence, and cognition.



Frieder Stolzenburg is professor of knowledge-based systems at the Harz University of Applied Sciences in Wernigerode and directs the Mobile Systems Laboratory at the Automation and Computer Science Department. His research interests include artificial intelligence, knowledge representation and logic, cognition and cognitive systems, and mobile robotics.



Karl-Heinz Weis is director of his company Weis Consulting e.K., founded in 2006. He permanently kept contact to the scientific community of artificial intelligence while working for the industry. His publications include works about case-based reasoning, experience management, machine learning, and neural networks.



Claus-Peter Wirth (Dr. rer. nat.) has a half-time position in the RatioLog project and is translator and editor in chief in the Hilbert-Bernays Project. He was a member of several SFBs, a SICSA distinguished visiting fellow, and a visiting senior research fellow at King's College London. His research interests include artificial intelligence and the history of logic and philosophy.

The final publication is available at:
<http://link.springer.com/article/10.1007/s13218-015-0377-9>